Automated State Space Abstraction for Reinforcement Learning

Team Members: Capt Robert Wright, Stephen Lin, Steven Loscalzo, AFRL/RISC
POC: Capt Robert Wright   (315) 330-2055   Robert.Wright@rl.af.mil

Key Concepts: Key Concepts: Reinforcement Learning, State Abstraction, State Aggregation, Feature Selection, Feature Extraction, Evolutionary Algorithms, and NeuroEvolution.

Summary: Reinforcement Learning (RL) is an attractive technology for enabling more autonomous systems that learn through direct interaction with their environment. However, RL technologies cannot be applied to most realistic domains without a state abstraction device to reduce the size and complexity of the sensed information within the environment. Current state abstraction methods are, in general, manually derived and fine tuned for each specific domain. They do not generalize well nor do they scale well as complexity increases.

In this project we are exploring automated methods for deriving state abstractions for RL algorithms. The research is focused on automating the process of state aggregation, feature selection, and feature extraction. Our approach is centered on the idea that more accurate abstractions will enable better reinforcement learning, thus the performance of an RL agent's learned policy is used as a metric for evaluating the quality of the state abstraction.

To date this research has produced several innovative algorithms designed to automate the state abstraction process. These algorithms have been evaluated extensively through empirical study and have we have demonstrated up to a fivefold improvement in learning speed on benchmark problems.